

**WHAT IS CLAIMED IS:**

5           1.       A method for copying a source data object to a destination data object while maintaining data coherency, the source data object being controlled by a source storage device controller, the destination data object being controlled by a destination storage device controller, comprising:

                  managing the source and destination storage device controllers using a  
10       remote application;

                  internally generating within the source storage device controller a snapshot version for each block of the source data object changed by one or more write operations to the block during the course of a copy operation; and

                  copying each block of the source data to a corresponding block in the  
15       destination data object in the absence of the snapshot version of the block and otherwise copying the snapshot version of the source data object block to the corresponding block in the destination data object, wherein data is directly transferred between the source and destination storage device controllers without traversing a file server for processing file requests, and wherein coherency of the data transferred between the source and  
20       destination storage device controllers is maintained without the file server maintaining or managing coherency maps.

                  2.       The method of claim 1, wherein each source data object block spans a byte  
25       range.

                  3.       The method of claim 1, further comprising instructing the source storage device controller to maintain an internal snapshot map to identify snapshot source data.

                  4.       The method of claim 3, further comprising performing a look-up in the  
30       source storage device controller's internal snapshot map to determine whether an existing write operation modifies a snapshot version of the source data object block.

5           5.       The method of claim 3, further comprising instructing the source storage device controller to update the source storage device controller's internal snapshot map after generating a snapshot of the source data.

10           6.       The method of claim 1, further comprising holding a write operation to the source data object until the source storage device controller's internal snapshot is updated.

15           7.       The method of claim 6, further comprising releasing the write operation to update the source data object if an internal snapshot version of the source data object block to be written to already exists.

            8.       An apparatus for copying a source data object distributed over one or more source controllers to a destination data object distributed over one or more destination controllers, comprising:

20                 a source storage controller to control access to the source data object, the source storage controller taking a snapshot version of each block in the source data object before updating any source data object block, the source storage controller taking the snapshot version without source data traversing a file server; and

25                 a replication manager to control multiple source storage device controllers, the replication manager enabling the one or more source controllers to take snapshots to provide a coherent copy of the destination data object.

30           9.       The apparatus of claim 8, wherein one or more commands are sent to the source storage controller.

            10.      The apparatus of claim 9, wherein one command is a copy command.

            11.      The apparatus of claim 9, wherein one command is a snapshot command.

5           12.     The apparatus of claim 9, wherein one or more of the commands are sent using an in-band protocol.

          13.     The apparatus of claim 9, wherein one or more of the commands are sent using an out-of-band protocol.

10           14.     A system for copying source data while maintaining data coherency, the source data having one or more blocks of data, the system comprising:

                  means for generating a snapshot version in a source data storage device, without source data traversing a file server, for each source data block changed by one or  
15       more write operations to the source data block during the course of a copy operation; and

                  means for copying, without the source data traversing the file server, each block of the source data to a corresponding block in the destination data in the absence of the snapshot version of the source data block and otherwise copying the snapshot version of the source data block to the corresponding block in the destination data.

20           15.     The system of claim 14, wherein each source data block spans a byte range.

          16.     The system of claim 14, further comprising a list of source data blocks to  
25       be copied which is reordered by the copying means to optimize copy speed.

          17.     The system of claim 14, wherein additional control data is inserted before and after the source data block while copying the source data block to the corresponding block in the destination data.

30           18.     The system of claim 16, wherein the list of blocks to be copied is buffered on the copying means while awaiting further copy requests.

          19.     The system of claim 14, wherein the block size is specified to the copying  
35       means so that fixed-size blocks are written to a destination controller device.

5           20.     The system of claim 14, further comprising means for maintaining a snapshot map to identify snapshot source data.

          21.     The system of claim 20, further comprising means for looking-up the snapshot map to determine whether an existing write operation modifies a snapshot  
10    version of the block.

          22.     A computer system for copying source data while maintaining data coherency, comprising:

                  a processor;

15               a source data storage device coupled to the processor to store source data, the source data having one or more blocks of data;

                  a destination data storage device coupled to the processor to store destination data;

                  means for generating a snapshot version in the source data storage device, without the source data traversing a file server, for each source data block changed by  
20    one or more write operations to the source data block during the course of a copy operation; and

                  means for copying, without the source data traversing the file server, each block of the source data to a corresponding block in the destination data in the absence of  
25    the snapshot version of the source data block and otherwise copying the snapshot version of the source data block to the corresponding block in the destination data.